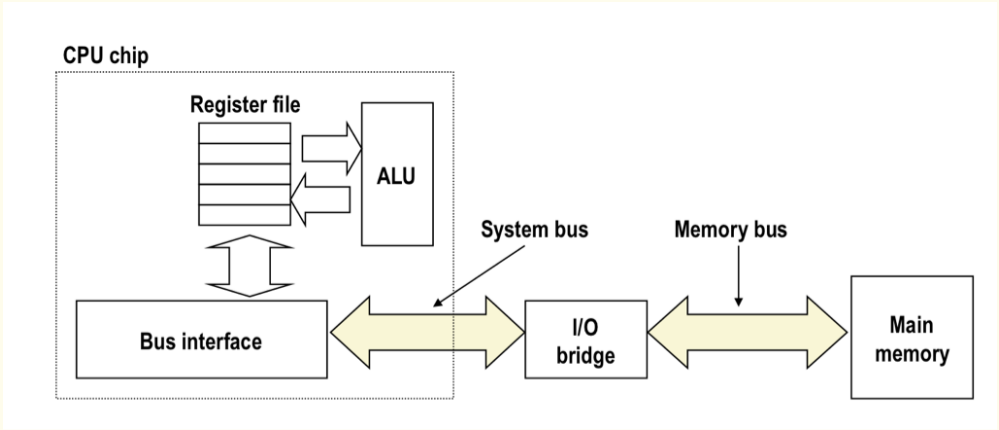


# Memory Hierarchy

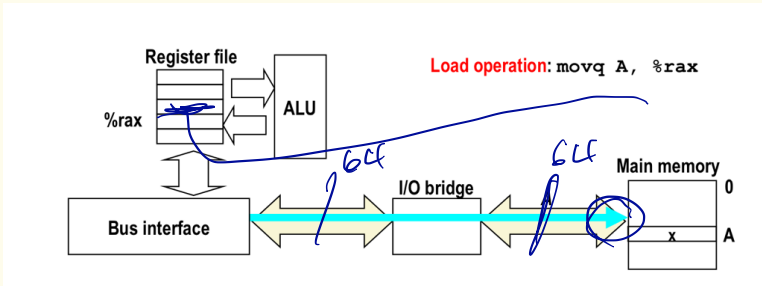


# Where Memory really fits in

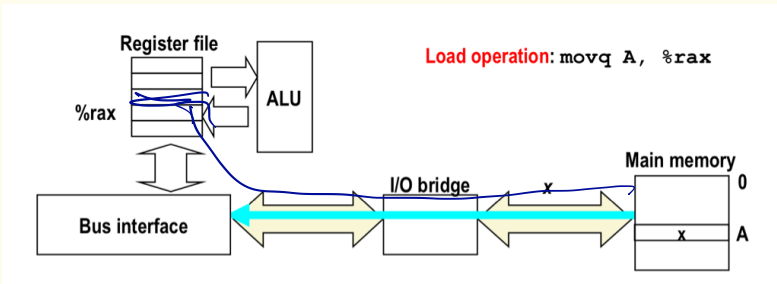


## Reading from memory

①




②




# Writing to Memory



①



②



## Disk Basics



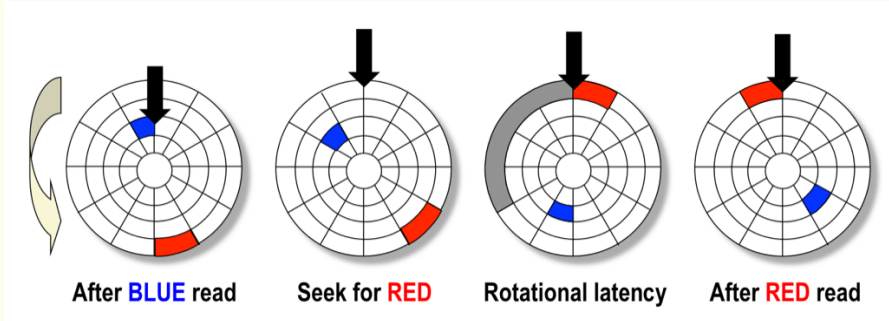
- Disks consist of **platters**, each with two **surfaces**.
- Each surface consists of concentric rings called **tracks**.
- Each track consists of **sectors** separated by **gaps**.

**Capacity = (# bytes/sector) x (avg. # sectors/track) x  
(# tracks/surface) x (# surfaces/platter) x  
(# platters/disk)**

**Example:**

- 512 bytes/sector
- 300 sectors/track (on average)
- 20,000 tracks/surface
- 2 surfaces/platter
- 5 platters/disk

**Capacity = 512 x 300 x 20000 x 2 x 5  
= 30,720,000,000  
= 30.72 GB**



$$T_{io} = T_{seek} + T_{rotate} + T_{transfer}$$

## ■ Given:

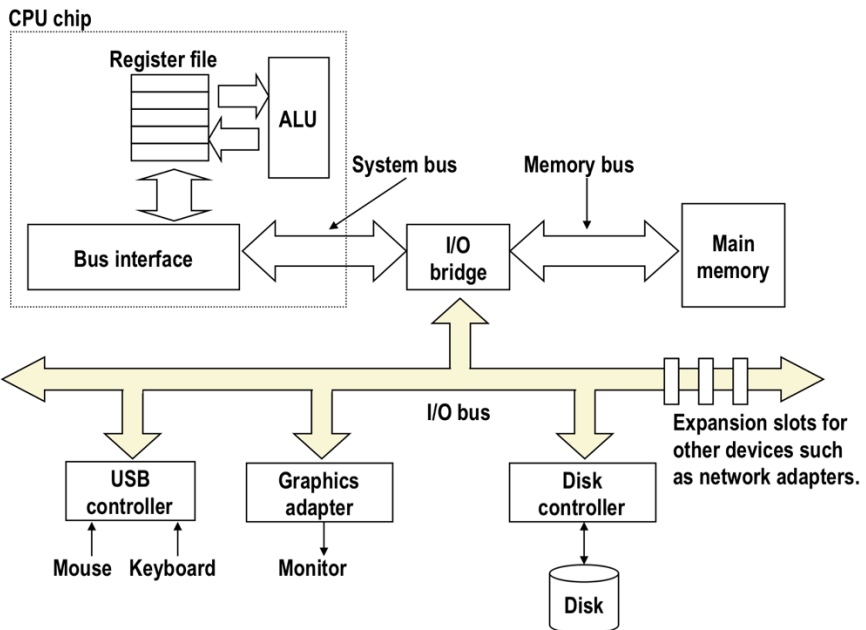
- Rotational rate = 7,200 RPM
- Average seek time = 9 ms.
- Avg # sectors/track = 400.

## ■ Derived:

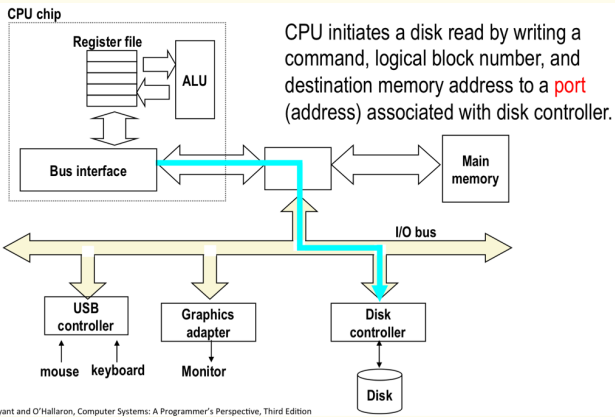
- Tavg rotation =  $1/2 \times (60 \text{ secs}/7200 \text{ RPM}) \times 1000 \text{ ms/sec} = 4 \text{ ms}$ .
- Tavg transfer =  $60/7200 \text{ RPM} \times 1/400 \text{ secs/track} \times 1000 \text{ ms/sec} = 0.02 \text{ ms}$
- Taccess = 9 ms + 4 ms + 0.02 ms

## ■ Important points:

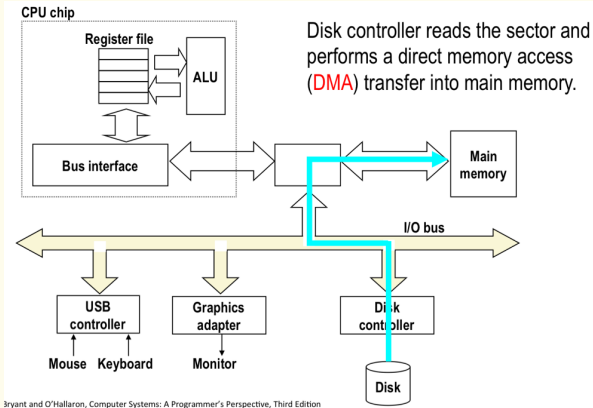
- Access time dominated by seek time and rotational latency.
- First bit in a sector is the most expensive, the rest are free.
- SRAM access time is about 4 ns/doubleword, DRAM about 60 ns
  - Disk is about 40,000 times slower than SRAM,
  - 2,500 times slower than DRAM.



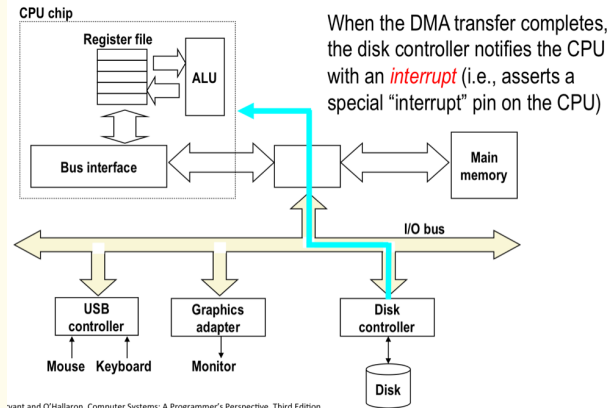
1

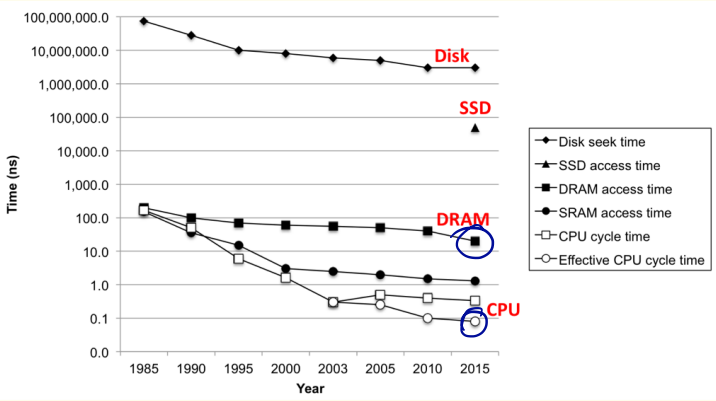


2



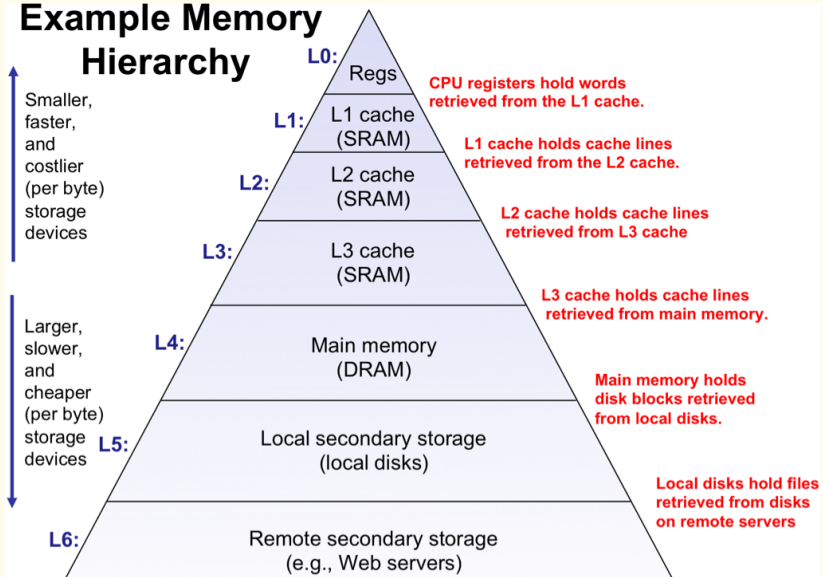
3





on to locality discussion

# Example Memory Hierarchy



- **Cache:** A smaller, faster storage device that acts as a staging area for a subset of the data in a larger, slower device.
- **Fundamental idea of a memory hierarchy:**
  - For each  $k$ , the faster, smaller device at level  $k$  serves as a cache for the larger, slower device at level  $k+1$ .
- **Why do memory hierarchies work?**
  - Because of locality, programs tend to access the data at level  $k$  more often than they access the data at level  $k+1$ .
  - Thus, the storage at level  $k+1$  can be slower, and thus larger and cheaper per bit.
- **Big Idea:** The memory hierarchy creates a large pool of storage that costs as much as the cheap storage near the bottom, but that serves data to programs at the rate of the fast storage near the top.



NB: Address is King

A. Interface to memory

CPU may deal in registers, but ultimately

movq <mem>, <reg>

movq <reg>, <mem>

yields an address to Bus

(and) Instruction<sub>0-a</sub> ← icount[PC]

locality: Defined relative to an address  
temporal  
spatial

Important for all to remember b/c

Programs with good locality run faster  
than programs with poor locality

and it is in programmer's control  
to favor locality

```
sum = 0;
for (i = 0; i < n; i++)
    sum += a[i];
return sum;
```

Which is better for locality, and why?

```
int sum_array_rows(int a[M][N])
{
    int i, j, sum = 0;


    for (i = 0; i < M; i++)
        for (j = 0; j < N; j++)
            sum += a[i][j];
    return sum;
}
```

```
int sum_array_cols(int a[M][N])
{
    int i, j, sum = 0;


    for (j = 0; j < N; j++)
        for (i = 0; i < M; i++)
            sum += a[i][j];
    return sum;
}
```

Does it really make a difference?

- **Cache memories** are small, fast SRAM-based memories managed automatically in hardware
  - Hold frequently accessed blocks of main memory
- **CPU looks first for data in cache**
- **Typical system structure:**



Cache




Smaller, faster, more expensive memory caches a subset of the blocks



Data is copied in block-sized transfer units


Memory




Larger, slower, cheaper memory viewed as partitioned into "blocks"

Request: 14

Cache




Memory




Request: 12

Cache




Memory




**Block b is stored in cache**

- **Placement policy:** determines where b goes
- **Replacement policy:** determines which block gets evicted (victim)



Define a cache organization by  $(S, E, B)$  + need  $m$ : number of address bits





Derived from above:

$M = 2^m$ : Maximum number of unique memory addresses

$s = \log_2(S)$ : Number of bits to represent set

$b = \log_2(B)$ : Number of bits to represent block offset


Address of word:




- 1 • **Locate set**
- 2 • **Check if any line in set has matching tag**
- 3 • **Yes + line valid: hit**
- 4 • **Locate data starting at offset**

Example 1:

Direct mapped: One line per set  
Assume: cache block size 8 bytes



②



③ + ④

